

REMARKS

I. INTRODUCTION

In response to the Office Action dated June 28, 2005, claims 1, 15, and 29 have been amended and claims 11, 25, and 39 have been cancelled. Claims 1-10 and 11-14 remain pending, claims 15-24, 26-38, and 40-42 have been withdrawn from consideration. Entry of these amendments, and re-consideration of the application, as amended, is requested.

II. PRIOR ART REJECTIONS

In paragraph (3) of the Office Action, claims 1-14 were rejected under 35 U.S.C. §102(e) as being anticipated by Adelman, U.S. Patent No. 6,006,259 (Adelman).

Applicants respectfully traverse these rejections.

Specifically, the independent claims and claim 11 where rejected as follows:

As to claim 1, Adelman teaches a method of providing a single system image in a clustered environment comprising:

- assigning an internet protocol (IP) address as a cluster IP address (see col.9, lines 22 – 30, Adelman discloses the assignment of IP address to a cluster commander);
- binding the cluster IP address to a node in a cluster (see col. 9, lines 22 – 30, Adelman discloses the binding of a cluster IP address to cluster members) ;
- receiving a client request directed to the cluster IP address(see col. lines 3 – 29, Adelman discloses messages received by the cluster);
- multicasting the request to all nodes in the cluster (see col. 9 lines 24 – 47 Adelman discloses multicasting Mac address to all the nodes in a cluster)
- each node in the cluster filtering the request based on a dynamically adjustable workload distribution function on each node, wherein the function is configured to allow a single node to process the client request (see col.9, Adelman discloses filtering the process by allowing cluster members to tell whether the incoming message must be processed by it) ;
- a single node in a cluster obtaining a response to the request (see col.9, Adelman discloses filtering the process by allowing cluster members to tell whether the incoming message must be processed by it) ;
- the single node inserting a cluster media access control (MAC) address into the response (see col. 9 lines 59 – col. 10 line 11 Adelman discloses MAC address of each cluster member) ;
- scnding the response from the single node to the client (see col. 8 lines 55 – 67 , Adelman discloses the notification that a request is accepted).

As to claim 11, Adelman the method of claim 1 wherein the workload distribution function distributes the workload by:

Representing a total workload observed by the cluster as a bitmap with a number of bits k (see col. 9, line 50-col. 10, line 11).

Obtaining a bit m by moding a source IP address of the client by the number of bits k (see col. 9, line 50-col. 10, line 11).

; and

assigning the client request to a cluster node that has a specified value at bit m (see col. 9, line 50-col. 10 line 11).

Applicants traverse the above rejections for one or more of the following reasons:

- (1) Adelman does not teach, disclose or suggest representing the total workload as a bitmap with a number of bits k ;
- (2) Adelman does not teach, disclose or suggest obtaining a bit m by performing a mod operation of a source IP address of the client by the number of bits k ; and
- (3) Adelman does not teach, disclose or suggest assigning a client request to a cluster node that has a specified value at a bit m .

As set forth in the claims as amended, a dynamically adjustable workload distribution function is on each node and allows a single node to process a client request. The distribution function distributes the workload using a specific methodology. Namely, the total workload observed is represented as a bitmap with a number of bits k . A bit m is obtained by MODing the source IP address of the client by the number of bits k . The client request is then assigned to a cluster node that has a specified value at the bit m .

In rejecting prior dependent claim 11 that contained the bitmap based workload distribution function, the Office Action relied on col. 9, line 50–col. 10, line 11 of Adelman which provides:

In the preferred embodiment, there is a mechanism for designating which cluster member is to process a message and allow the other members to disregard the message without inadvertently sending a "reset" message to the originating client. The preferred embodiment makes use of a "filter" process in each cluster member which calculates a hash function using certain fields of the incoming message header. This hash calculation serves as a means of both assigning a work unit number to a message and assigning a work unit to a particular cluster member for processing. This technique allows a cluster member to tell whether the incoming message must be processed by it, therefore the possibility of an inadvertent "reset" message is precluded. It is noted that other solutions to this problem of "how to get the work to the right member of the cluster with minimum overhead" could include a hardware filter device sitting between the network and the cluster wherein the hardware filter would do the member assignment and load balancing function. Note that since all cluster members have the same MAC address, all cluster members get all messages and the way they tell whether they must process the message further is to calculate the work unit number using the hashing method shown above and then to check the resulting work unit number against their work load table to see if it is assigned to them. If not they dump the message from their memory. This is a fast and efficient scheme for dumping messages that the units need not process further and yet it provides an efficient basis for load-balancing and efficient fail-over handling when a cluster member fails.

As can be seen from the above text, there is not even a remote suggestion to use a bitmap with a particular number of bits or performing a particular bit in the bitmap by performing a MOD operation using a source IP address and a number of bits. Instead, the cited portion merely provides that to tell whether a particular member should process a message, a work unit number is calculated using a hashing method. The hashing method is clearly set forth in Adelman at col. 10, line 21 – col.

11, line 46. Namely, the IP source and destination addresses are hashed modulo the number 1024 to produce an index value (see col. 10, lines 25-27). More specifically, the field containing the IP addresses, IP protocol, TCP/UDP port number, and session and tunnel ID fields are added together (using a logical XOR) and shifted to produce a number between 0 and 1023 (see col. 10, lines 31-35). The value is then used as a pointer to a work set to see if the work set has been assigned to a particular member (see col. 11, lines 17-20)

Such a hash calculation does not even remotely teach, describe, suggest, or allude to representing a total workload observed by a cluster as a bitmap with a particular number of bits. Further, the process of obtaining a particular bit by MODing a source IP address by the number of bits in the bitmap and then assigning a client request to a node with a particular value at the obtained bit is not even remotely hinted at.

In fact, Applicants submit that the above portions of Adelman merely describe how a member determines whether the incoming client request should be processed by them or not. Adelman's workset or the creation of the workset would establish how work is assigned to a particular client. Adelman describes the work assignment table that contains work-unit hash number and the cluster member id assigned to a work-unit in col. 6, lines 20-32. The specific assignment of workload is described in col. 8, lines 31-49 as follows:

As indicated above, the master periodically sends out a master keepalive message containing the cluster member list, the adaptive keepalive interval (which is described in more detail below) and the current set of work assignments for each member which is used only for diagnostic purposes. (See FIG. 8C). In addition, the master periodically (in the preferred embodiment every 2 seconds) checks the load-balance of the cluster members. In FIG. 8D when the timer expires 855 the master calculates the load difference between most loaded (say "K") and least loaded (say "J") cluster member 857 and then asks "would moving 1 work unit from most loaded (K) to least loaded (J) have any effect?" that is, if $K > J$ is $K - 1 \geq J + 1$? 859. If so then the master sends a "work de-assign" request to the most loaded member with the least loaded member as the target recipient 863 and then the master checks the load numbers again 865. If the result of moving 1 work unit would not leave the least loaded less than or equal to the most loaded 860 then the master makes no reassignments and exits 861.

As described in this text, a load difference is calculated and a question is asked regarding whether moving work from one unit to another unit would have an effect and adjusting the workload based on the comparison. However, such a teaching does not remotely refer or allude to a bitmap, a MOD operation, or looking at a particular bit for a specific value.

In fact throughout Adelman, the specific limitations of the amended independent claims are not even hinted at. Further, the various elements of Applicants' claimed invention together provide

operational advantages over the systems disclosed in Adelman. In addition, Applicants' invention solves problems not recognized by Adelman.

Thus, Applicants submit that independent claim 1 is allowable over Adelman. Further, dependent claims 2-14 are submitted to be allowable over Adelman in the same manner, because they are dependent on independent claim 1, and because they contain all the limitations of the independent claims. In addition, dependent claims 2-14 recite additional novel elements not shown by Adelman.

In addition, Applicants again reassert the traversal of the restriction requirement in that the steps performed in the method claim are clearly similar to those performed by the master node and storage cluster of the non-elected claims. Further, the non-elected claims are merely apparatus and article of manufacture claims with almost identical limitations to that of the method claim. Accordingly, Applicants again respectfully request withdrawal of the restriction requirement.

III. CONCLUSION

In view of the above, it is submitted that this application is now in good order for allowance and such allowance is respectfully solicited. Should the Examiner believe minor matters still remain that can be resolved in a telephone interview, the Examiner is urged to call Applicants' undersigned attorney.

Respectfully submitted,

Ying Chen et al.

By their attorneys,

GATES & COOPER LLP

Howard Hughes Center
6701 Center Drive West, Suite 1050
Los Angeles, California 90045
(310) 641-8797

Date: September 26, 2005

JSF/bjs

By: 

Name: Jason S. Feldman

Reg. No.: 39,187